

Energy Management Policies for Energy-Neutral Source-Channel Coding

P. Castiglione*, O. Simeone[†], E. Erkip[‡], and T. Zemen*

*Forschungszentrum Telekommunikation Wien, Austria

[†]CWCSPR, ECE Dept, NJIT, New Jersey, USA

[‡]Dept. of ECE, Polytechnic Inst. of NYU, New York, USA

Abstract

In cyber-physical systems where sensors measure the temporal evolution of a given phenomenon of interest and radio communication takes place over short distances, the energy spent for source acquisition and compression may be comparable with that used for transmission. Additionally, in order to avoid limited lifetime issues, sensors may be powered via energy harvesting and thus collect all the energy they need from the environment. This work addresses the problem of energy allocation over source acquisition/compression and transmission for energy-harvesting sensors. At first, focusing on a single-sensor, energy management policies are identified that guarantee a maximal average distortion while at the same time ensuring the stability of the queue connecting source and channel encoders. It is shown that the identified class of policies is optimal in the sense that it stabilizes the queue whenever this is feasible by any other technique that satisfies the same average distortion constraint. Moreover, this class of policies performs an independent resource optimization for the source and channel encoders. Analog transmission techniques as well as suboptimal strategies that do not use the energy buffer (battery) or use it only for adapting either source or channel encoder energy allocation are also studied for performance comparison. The problem of optimizing the desired trade-off between average distortion and delay is then formulated and solved via dynamic programming tools. Finally, a system with multiple sensors is

This work has been submitted to the IEEE for publication. Copyright may be transferred without notice, after which this version may no longer be accessible. Part of this paper has been accepted for publication at the 9th Int. Symposium on Modeling and Optimization in Mobile, Ad Hoc, and Wireless Networks (WiOpt 2011), Princeton, New Jersey, USA on May 2011. This work was supported by the Austria Science Fund (FWF) through grant NFN SISE (S106). The Telecommunications Research Center Vienna (FTW) is supported by the Austrian Government and the City of Vienna within the competence center program COMET.

considered and time-division scheduling strategies are derived that are able to maintain the stability of all data queues and to meet the average distortion constraints at all sensors whenever it is feasible.

I. INTRODUCTION

In the “smart world”, wireless sensor networks (WSNs) play a central role in bridging the real and the digital worlds [1]. WSNs are typically designed under the assumptions that communication resources are limited by the energy available in the battery and that the most significant source of energy expenditure is radio transmission. However, modern cyber-physical systems are expected to operate over a virtually infinite lifetime. This can only be achieved by overcoming the limitations of battery-powered sensors and allowing the sensors to *harvest* the energy needed for their operation from the environment, e.g., in the form of solar, vibrational or radio energy [2], [3]. The regime of operation in which the system operates in a fully self-powered fashion is referred to as *energy neutral* [4]. Moreover, when sensors are tasked with acquiring complex measures, such as long time sequences of given phenomena of interest, and when transmission takes place over small distances, the energy cost of running the source acquisition system (sensing, sampling, compression) may be comparable with that of radio transmission [5], [6].

Based on the discussion above, in this paper, we address the problem of energy management for a WSN in which sensors are powered via energy harvesting and in which source acquisition and radio transmission have comparable energy requirements. We first focus on a system with a single sensor communicating to a single receiver, as shown in Fig. 1, in order to concentrate on the main aspects of the problem. The sensor is equipped with a battery in which the harvested energy is stored. In each time slot, the sensor acquires a time sequence for the phenomenon of interest, which is characterized by a measurement signal-to-noise ratio (SNR) and autocorrelation, and stores the resulting bits, after possible compression, into a data queue. At the same time, it transmits a number of bits from the data queue to the fusion center over a fading channel with an instantaneous channel SNR. Based on the statistics of the energy harvesting process, and based on the current states of the measurement quality, of channel SNR, and of the data queue, the energy management unit must perform energy allocation between source acquisition

and data transmission so as to optimally balance competing requirements such as distortion of the reconstruction at the receiver, queue stability and delay. This optimization problem is the main subject of this work. We further extend our analysis to the problem of scheduling multiple sensors that are communicating to the same receiver.

The model at hand is inspired by the work in [5], [6] and [7]. In [7], the energy-harvesting sensor allocates power to data transmission over different channel SNRs, since the bit arrival process is assumed to be given and not subject to optimization. This is unlike our work in which a key problem is that of allocating resources between transmission and source compression in order to guarantee given constraints such as distortion and queue stability. The problem of energy allocation between source compression and transmission was instead first studied in [5], [6], but in power-limited systems with no energy-harvesting capabilities.

Other related works pertain to the study of energy-harvesting WSNs. This is a growing field with recent significant contributions. Here we only point to the works that are most related to ours, besides the ones already mentioned above. An information-theoretic analysis of a single-sensor system with energy-harvesting is presented in [8], [9], where it is shown that energy-harvesting does not affect the capacity of the channel, as long as one assumes that the battery has an arbitrarily large storage capacity. An optimal strategy for a single-sensor system that can control both the “acceptance rate” of the arriving bits and the power allocation with the aim of maximizing the throughput under stability constraints is developed in [10]. Optimal scheduling is instead studied in [11], [12], [13]. The effect of a finite battery is studied in [14], where the trade-off between achievable rate and battery discharge probability is characterized. It is noted that all these works do not model the aspect of source acquisition and processing.

The main contributions of this work are summarized as follows. (i) We propose a simple, but general, model for an energy-harvesting sensor operating over a time-varying channel (Sec. II). (ii) For a single-sensor system, we design a novel class of *distortion-optimal energy-neutral* resource allocation policies that are able to stabilize the data queue and, simultaneously, to meet an average distortion constraint, whenever it is feasible by any policy (Sec. III-A). For the case where multiple sensors access the same uplink channel in time division, we identify a distortion-

optimal energy-neutral class of *scheduling* policies (Sec. V). (iii) We compare the performance of the optimal policies with a number of less complex strategies, such as “analog” techniques [15] (Sec. III-C) and fixed time division multiple access (TDMA) scheduling strategies (Sec. V-A). (iv) Finally, we formulate the problem of optimizing a desired trade-off between average delay and distortion, which is solved via dynamic programming tools (Sec. IV).

II. SYSTEM MODEL

In this section, we introduce the system model, main assumptions and problem definition.

We consider a system in which a single sensor communicates with a single receiver as depicted in Fig. 1. The extension of this system to the case of multiple sensors and a single receiver is studied in Sec. V. In most of the paper, we assume that the sensor performs separate source and channel coding, as described in the following. A different approach is considered in Sec. III-C.

Time is slotted. The energy $E_k \in \mathbb{R}_+$ harvested in time-slot k is stored in an “energy buffer”, also referred to as battery, with infinite size. For convenience, the energy E_k is normalized to the number N of channel discrete-time symbols available for communication in each time slot, also referred to as channel uses. The energy arrival E_k is assumed to be a stationary ergodic process. The probability density function (pdf) of E_k is $p_E(e)$. The energy \tilde{E}_{k+1} available for use at slot $k+1$ is the residual energy from the previous slot plus the energy arrival at time-slot $k+1$. This evolves as

$$\tilde{E}_{k+1} = \left[\tilde{E}_k - (T_{s,k} + T_{t,k}) \right]^+ + E_{k+1}, \quad (1)$$

where $T_{s,k}$ and $T_{t,k}$ account for the energy spent in slot k per channel use for source acquisition and data transmission, respectively, as discussed below. Notice that the energy arriving at time slot $k+1$ is immediately available for use in that slot.

The sensor measures M samples of a given source during each slot. The quality of such observation in slot k depends on a parameter $Q_k \in \mathcal{Q}$, which is assumed to be a stationary ergodic process over the time slots k . For instance, the sensor may perform measurements of the phenomenon of interest whose SNR Q_k changes across blocks k due to source movement or environmental factors affecting the measurement quality. The set \mathcal{Q} is assumed to be discrete

and finite, and the (stationary) probability mass function (pmf) for Q_k is given by $\Pr(q) = \Pr(Q_k = q)$, for $q \in \mathcal{Q}$. The sensor acquires the source in a lossy fashion. The loss, due to sampling, analog-to-digital conversion and compression, is characterized by distortion $D_k \in \mathbb{R}^+$, as measured with respect to some distortion metric such as the mean square error (MSE).

The number of bits generated by the source encoder at the sensor at slot k is $X_k = f(D_k, T_{s,k}, Q_k)$, where f is a given function of the distortion level D_k , of the energy per channel use allocated to the source encoder $T_{s,k}$ and on the observation state Q_k . The resulting bit stream is buffered in a first-input-first-output (FIFO) data queue with queue length \tilde{X}_k . The function $f(D_k, T_{s,k}, Q_k)$ is assumed to be separately continuous convex and non-increasing in D_k and $T_{s,k}$. For simplicity, we will denote such functions also as $f^q(D_k, T_{s,k}) = f(D_k, T_{s,k}, Q_k = q)$. Some examples for function f will be provided below in Sec. II-A.

The fading channel between sensor and destination is characterized by a process H_k , assumed to be stationary ergodic, where $H_k \in \mathcal{H}$, with set \mathcal{H} being discrete and finite in order to ease the numerical evaluations. We assume a slowly time-variant scenario. The pmf of H_k is given by $\Pr(h) = \Pr(H_k = h)$, for $h \in \mathcal{H}$. The *channel encoder* uses the channel N times per slot, and the transmission requires $T_{t,k}$ energy per channel use. A maximum number $g(H_k, T_{t,k})$ of bits per slot can be delivered successfully to the destination. The channel rate function $g(H_k, T_{t,k})$ is assumed to be continuous, concave, non-decreasing in $T_{t,k}$, and $g(H_k, 0) = 0$. We also use the notation $g^h(T_{t,k}) = g(H_k = h, T_{t,k})$. An example is the Shannon capacity on the complex additive white Gaussian noise (AWGN) channel $g^h(T_{t,k}) = N \times \log(1 + hT_{t,k})$ [16]. We remark that adopting the Shannon capacity implies the use of rate-adaptive schemes with sufficiently long codewords so that the block error probability becomes negligible. For the given function $g(H_k, T_{t,k})$, the channel encoder takes $\min[\tilde{X}_k, g(H_k, T_{t,k})]$ bits from the data buffer, using the selected transmission energy $T_{t,k}$. Note that we do not consider the effects of channel errors nor the costs of channel encoding/decoding and of channel state information feedback, which are beyond the scope of the present paper and subject to future work.

Based on the discussion above, the data queue evolves as

$$\tilde{X}_{k+1} = \left[\tilde{X}_k - g(H_k, T_{t,k}) \right]^+ + f(D_k, T_{s,k}, Q_k). \quad (2)$$

To illustrate the trade-offs involved in the energy allocation between $T_{t,k}$ and $T_{s,k}$, we remark that, by providing more energy $T_{s,k}$ to the source encoder, one is able, for the same distortion level D_k , to reduce the number $f(D_k, T_{s,k}, Q_k)$ of bits to be stored the data buffer. At the same time, less energy $T_{t,k}$ is left for transmission, so that the data buffer is emptied at a lower rate $g(H_k, T_{t,k})$. Viceversa, one could use less energy to the source encoder, thus producing more bits $f(D_k, T_{s,k}, Q_k)$, so that more energy would be available to empty the data buffer.

A. Rate-Distortion-Energy Trade-Off

In the following, we present some examples for function $f^q(D_k, T_{s,k})$, as available in the literature. Recall that this function provides the trade-off between the distortion D_k , the energy consumption $T_{s,k}$ and the number of bits produced by the source encoder.

Example 1. Consider the observation model $R_{k,i} = \sqrt{Q_k}U_{k,i} + Z_{k,i}$, where M samples of the random process $U_{k,i}$, for $i \in \{1, \dots, M\}$, are measured during the slot k and each measurement $R_{k,i}$ is affected by Additive White Gaussian Noise (AWGN) Z_k with unitary variance. Parameter Q_k represents the observation SNR in slot k . From [5], an approximated and analytically tractable model for $f^q(D_k, T_{s,k})$ is

$$f^q(D_k, T_{s,k}) = \frac{N}{b} \times f_1^q(D_k) \times f_2(T_{s,k}), \quad (3)$$

where $b = N/M$ is the bandwidth ratio and $f_2(T_{s,k}) = \zeta \times \max[(bT_{s,k}/T_s^{max})^{-1/\eta}, 1]$ models the rate-energy trade-off at the source encoder. The parameter $\zeta > 1$ is related to the efficiency of the encoder, the coefficient $1 \leq \eta \leq 3$ is specified by the the given processor [17] and parameter T_s^{max} upper bounds the energy $T_{s,k}$ that can be used by the source encoder. Function $f_1^q(D_k)$ is a classical rate-distortion function [16]. For the model described in this example, assuming that the source is independent identically distributed (i.i.d) in time with $U_{k,i} \sim \mathcal{N}(0, d_{max})$, the

rate-distortion trade-off is given by

$$f_1^q(D_k) = \left(\log \frac{d_{max} - d_{mmse}}{D_k - d_{mmse}} \right)^+, \text{ where } d_{mmse} = \left(\frac{1}{d_{max}} + q \right)^{-1}, \quad (4)$$

where d_{mmse} is the estimation minimum MSE (MMSE) for the estimate of $U_{k,i}$ given $R_{k,i}$ [18]. Notice that the distortion D_k is upper bounded by d_{max} and lower bounded by d_{mmse} .

Example 2. The sensor observes M samples of a first-order Gaussian Markov source $[U_{k,1}, U_{k,2}, \dots, U_{k,M}] \in \mathbb{R}^M$ with correlation function given by $\mathbb{E}[U_k U_{k+j}] = d_{max} Q_k^{|j|}$, where parameter $0 \leq Q_k \leq 1$ is the correlation coefficient between the samples measured in slot k . Notice that the larger Q_k is, the easier it is for the compressor to reduce the bit rate for a given distortion due to the increased correlation. Adapting results from [6], if the source encoder uses a transform encoder [19], the optimal compressor produces a number of bits equal to

$$f^q(D_k, T_{s,k}) = \frac{N}{b} \times [f_1(D_k) + f_2^q(T_{s,k})]^+, \quad (5)$$

with $f_1(D_k) = \log \frac{\zeta d_{max}}{D_k}$, where parameter $\zeta \geq 1$ depends on the type of quantizer, and

$$f_2^q(T_{s,k}) = \log(1 - q^2) \times \frac{T_{s,k} - \nu/b}{T_{s,k}}, \quad (6)$$

with given parameter ν , which sets a lower bound on the energy $T_{s,k}$ as $T_{s,k} \geq \nu/b$. Equations (5)-(6) are obtained by assuming, similar to [6], that the energy required for source compression is proportional to the size of the transform encoder. Finally, notice that, since the compression rate must be positive, D_k is upper bounded by $\zeta d_{max} (1 - q^2)^{\frac{T_{s,k} - \nu/b}{T_{s,k}}}$.

B. Problem Definition

At each time slot k , a *resource manager* must determine the distortion D_k and the energies $T_{s,k}$ and $T_{t,k}$ to be allocated to the source and channel encoder, respectively. The decision is taken according to a policy $\pi := \{\pi_k\}_{k \geq 1}$, where $\pi_k := \{D_k(S^k), T_{s,k}(S^k), T_{t,k}(S^k)\}$ determines parameters $(D_k, T_{s,k}, T_{t,k})$ as a function of the present and past states $S^k = \{S_1, \dots, S_k\}$ of the system, where the $S_i = \{\tilde{E}_i, \tilde{X}_i, Q_i, H_i\}$ accounts for the state of the available energy \tilde{E}_i , for

the data buffer \tilde{X}_i , for the the source observation state Q_i and the channel state H_i . We define the set of all policies as Π . Policies can be optimized according to different criteria. In Sec. III we adopt stability under an average distortion criterion as criterion of interest, while Sec. IV addresses the optimization of the trade-off between distortion and delay.

III. STABILITY UNDER A DISTORTION CONSTRAINT

In this section, we adopt as performance criterion the stability of the data queue connecting source and channel encoders. We also impose the constraint that the policy guarantees the following condition on the long-term distortion:

$$\lim_{n \rightarrow \infty} \frac{1}{n} \sum_{k=1}^n D_k \leq \bar{D} \quad (7)$$

for a fixed maximum average distortion level \bar{D} tolerated by the system. We define a policy as \bar{D} -feasible if it guarantees the stability of the data queue connecting source and channel encoders under the average distortion constraint (7). Recall that stability of the data queue holds if the distribution of \tilde{X}_k is asymptotically *stationary* and *proper*, i.e., $\Pr(\tilde{X}_k = \infty) \rightarrow 0$ [20].

A. Distortion-Optimal Energy-Neutral Class of Policies

For a given distortion \bar{D} , our objective in this section is to identify a class of policies that is able to stabilize the data queue and satisfy the distortion constraint (7) as long as this is possible. We refer to this class of policies as *distortion-optimal energy-neutral*. Notice that this definition generalizes that of “throughput optimal” policies [21] considered in related works such as [7], where only the stability constraint is imposed. By definition, a *distortion-optimal energy-neutral class* of policies $\Pi^{do} \subseteq \Pi$ contains at least one \bar{D} -feasible policy. For instance, the set Π of all policies is clearly distortion-optimal energy-neutral. However, this is a rather unsatisfying solution to the problem. In fact, it does not help in any way to identify a \bar{D} -feasible policy for a given system setup. Instead, we want to identify a smaller class Π^{do} , which is parametrized in a way that makes it easy to evaluate a \bar{D} -feasible policy. The propositions below identify a

distortion-optimal energy-neutral class of policies for the separate source and channel encoders model depicted in Fig. 1 and described in Sec. II.

Proposition 3. *For a given distortion \bar{D} , a necessary condition for the existence of a \bar{D} -feasible policy is the existence of a set of parameters $D^q \geq 0$, $T_s^q \geq 0$ for $q \in \mathcal{Q}$, $T_t^h \geq 0$ for $h \in \mathcal{H}$, and $0 < \alpha < 1$ such that*

$$\sum_q \Pr(q) f^q(D^q, T_s^q) < \sum_h \Pr(h) g^h(T_t^h), \quad \sum_q \Pr(q) D^q \leq \bar{D}, \quad (8)$$

$$\sum_q \Pr(q) T_s^q \leq (1 - \alpha) \mathbb{E}[E_k], \quad \text{and} \quad \sum_h \Pr(h) T_t^h \leq \alpha \mathbb{E}[E_k]. \quad (9)$$

Remark 4. Parameters D^q , T_s^q , T_t^h and α , whose existence is necessary for the existence of a \bar{D} -feasible policy according to Proposition 3, have a simple interpretation. In particular, T_s^q , D^q can be read as the average energy and distortion that the source encoder selects when the observation state is $Q_k = q$, whereas T_t^h can be seen as the average energy that channel encoder draws from the available energy for transmission when the channel state is $H_k = h$. Moreover, condition (8)-left is necessary for the stability of the data queue, condition (8)-right is necessary to satisfy the constraint (7), and conditions (9) are necessary for energy neutrality. This interpretation will be used below to derive a class of distortion-optimal energy-neutral policies.

Proof: The processes $f(D_k, T_{s,k}, Q_k)$ and $g(H_k, T_{t,k})$ must be asymptotically stationary ergodic for queue (2) to be asymptotically stationary. Hence, the policy π must be asymptotically stationary. Under this assumption, the necessary condition for the distribution of \tilde{X}_k to be asymptotically proper is $\mathbb{E}_\pi[f(D_k, T_{s,k}, Q_k)] < \mathbb{E}_\pi[g(H_k, T_{t,k})]$ from standard results on G/G/1 queues (see any reference on queuing theory, e.g., [20, Ch.3]). Notice that the average \mathbb{E}_π , with respect to the joint distribution of the state variables E_k, Q_k, H_k , is explicitly dependent on the policy π_k . From this condition, since f is separately convex in $D_k, T_{s,k}$ and g is concave in $T_{t,k}$, we have the following necessary condition $\sum_q \Pr(q) f^q(\mathbb{E}_\pi[D_k | Q_k = q], \mathbb{E}_\pi[T_{s,k} | Q_k = q]) < \sum_h \Pr(h) g^h(\mathbb{E}_\pi[T_{t,k} | H_k = h])$, where we have used Jensen inequality on both sides. Defining $D^q = \mathbb{E}_\pi[D_k | Q_k = q]$, $T_s^q = \mathbb{E}_\pi[T_{s,k} | Q_k = q]$, and $T_t^h = \mathbb{E}_\pi[T_{t,k} | H_k = h]$, the condition

(8)-left is then proved. As for (9), we consider that, from (1), we must have $\frac{1}{K} \sum_{k=1}^K (T_{s,k} + T_{t,k}) \leq \frac{1}{K} \sum_{k=1}^K E_k + \frac{\tilde{E}_0}{K}$, for $K \geq 1$, and the initial state of the energy buffer \tilde{E}_0 . Then, for a stationary ergodic policy π , we get $\mathbb{E}_\pi [T_{s,k}] + \mathbb{E}_\pi [T_{t,k}] \leq \mathbb{E} [E_k]$, where $\frac{1}{K} \sum_{k=1}^K T_{s,k} \rightarrow \mathbb{E}_\pi [T_{s,k}]$, $\frac{1}{K} \sum_{k=1}^K T_{t,k} \rightarrow \mathbb{E}_\pi [T_{t,k}]$, and $\frac{1}{K} \sum_{k=1}^K E_k + \frac{\tilde{E}_0}{K} \rightarrow \mathbb{E} [E_k]$. Given the definitions and the inequality above, (9) are proved, having set $\alpha = \mathbb{E}_\pi [T_{t,k}] / \mathbb{E} [E_k]$. To conclude, for (8)-right, we observe that the distortion constraint (7) is satisfied. \blacksquare

We now look for a distortion-optimal energy-neutral class of policies. To this end, based on Proposition 3, it is enough to exhibit a class of policies such that it contains a \bar{D} -feasible policy as long as the necessary conditions (8)-(9) are satisfied for some set of parameters D^q , T_s^q , T_t^h and α . Proposition 3 suggests that it is possible to find \bar{D} -feasible policies that select D_k and $T_{s,k}$ based on the observation state Q_k only, whereas the selection of $T_{t,k}$ depends on the channel state H_k only. Based on this consideration, let us define the class of policies Π^{do}

$$\Pi^{do} \begin{cases} D_k = D^q, T_{s,k} = \min \left[(1 - \alpha) \tilde{E}_k - \epsilon, T_s^q \right] & \text{for } Q_k = q \\ T_{t,k} = \min \left[\alpha \tilde{E}_k - \epsilon, T_t^h \right] & \text{for } H_k = h \end{cases} \quad (10)$$

where $D^q \geq 0$, $T_s^q \geq 0$ for $q \in \mathcal{Q}$, $T_t^h \geq 0$ for $h \in \mathcal{H}$, and $0 < \alpha < 1$ are fixed design parameters.

Proposition 5. *A policy in Π^{do} is \bar{D} -feasible if conditions (8) hold, along with*

$$\sum_q \Pr(q) T_s^q \leq (1 - \alpha) \mathbb{E} [E_k] - \epsilon, \quad \text{and} \quad \sum_q \Pr(q) T_t^h \leq \alpha \mathbb{E} [E_k] - \epsilon. \quad (11)$$

Remark 6. The sufficient conditions in Proposition 5 for the policies in Π^{do} to be \bar{D} -feasible coincide, for $\epsilon \rightarrow 0$, with the necessary conditions derived in Proposition 3. Therefore Π^{do} contains a \bar{D} -feasible policy any time the necessary conditions of Proposition 3 hold. As discussed above, this implies that the set Π^{do} is a distortion-optimal energy-neutral class. Moreover, it should be noted that the class Π^{do} , given (10), is parametrized by a small number of parameters and the policies in Π^{do} perform separate resource allocation optimizations for the source and channel encoders. In particular, the energy allocated to the source encoder $T_{s,k}$ only depends on

the observation state Q_k , and not on the channel quality H_k , whereas the energy $T_{t,k}$ for the channel encoder only depends on H_k , and not on Q_k . The energy allocation between the two encoders is governed by a single parameter $0 < \alpha < 1$. This entails that, once this parameter is fixed, and thus the energy budget available at the two encoders is fixed, resource allocation at the two encoders can be done separately without loss of optimality.

Proof: For $0 < \alpha < 1$ such that $\sum_q \Pr(q) T_s^q \leq (1 - \alpha) \mathbb{E}[E_k] - \epsilon$ and $\sum_q \Pr(q) T_t^h \leq \alpha \mathbb{E}[E_k] - \epsilon$, with ϵ small, we obtain that $\Pr(\tilde{E}_k = \infty) = 1$ asymptotically. This is true since $\mathbb{E}[T_{s,k} + T_{t,k}] < \mathbb{E}[E_k]$ in the system (1), so that the energy harvested is larger than the energy consumed on average and the energy queue is not stable [20, Ch.3] (see also [7] for the same argument). This leads to the asymptotically infinite size of the stored energy, as the buffer capacity is assumed infinite. Therefore we have, $T_{s,k}(Q_k = q) \rightarrow T_s^q$ and $T_{t,k}(H_k = h) \rightarrow T_t^h$ from (10). Notice that this argument shows that in (10) one can substitute α for any number between 0 and 1 leading to the same sufficient conditions (8) and (11). Due to the stationarity and ergodicity of the processes Q_k and H_k , $f(D_k, T_{s,k}, Q_k)$ and $g(H_k, T_{t,k})$ are stationary ergodic, and $\sum_q \Pr(q) f^q(D^q, T_s^q) < \sum_h \Pr(h) g^h(T_t^h)$ is the sufficient condition for the stability of the queue \tilde{X}_k [20, Ch.3]. Being $\lim_{n \rightarrow \infty} \frac{1}{n} \sum_{k=1}^n D_k = \sum_q \Pr(q) D^q$, for D^q such that $\sum_q \Pr(q) D^q \leq \bar{D}$ the class of policies Π^{do} satisfies the constraint (7). ■

Remark 7. A problem of interest is to find the minimal distortion \bar{D} for which the set of distortion-optimal energy-neutral policies Π^{do} is not empty. In other words, assessing the minimal distortion that can be supported without causing the data queue to be unstable. Given the separate nature of the source and channel energy allocations, it can be seen that one should optimize both terms in (8)-left separately, once the optimal value for α has been found. In particular, when $g(H_k, T_{t,k})$ is the Shannon capacity, the policy $T_{t,k}$ that minimizes \bar{D} is the *water-filling* [16].

B. Suboptimal Classes of Policies

In Sec. III-A, a distortion-optimal energy-neutral class Π^{do} has been identified. This class of policies, as made clear by the proof of Proposition 5 requires infinite energy storage capabilities at the sensor node. Let us instead consider the class of greedy policies Π^{sub1} that do not use the

energy buffer but allocates all the energy arrival E_k to source and channel coding according to a fraction $0 \leq \alpha^{q,h} \leq 1$ that depends on both source $Q_k = q$ and channel $H_k = h$ states:

$$\Pi^{sub1} \left\{ D_k = D^{q,h}, T_{s,k} = \alpha^{q,h} E_k, T_{t,k} = (1 - \alpha^{q,h}) E_k \quad \text{for } Q_k = q \text{ and } H_k = h \right. \quad (12)$$

where the distortion $D^{q,h} \geq 0$ also depends on both source and channel states. Notice that this is unlike the class of distortion-optimal energy-neutral policies (10) in which, as explained in Remark 6, energy allocation is done independently for source (only based on Q_k) and channel decoder (only based on H_k). Here, parameters $\alpha^{q,h}, D^{q,h}$ are selected on the basis of both channel and source states Q_k and H_k to partially compensate for the loss due to the greedy approach. For further reference, we also consider the subclass of policies $\Pi^{sub2} := \Pi^{sub1}|_{\alpha^{q,h}=\alpha}$ for all (q,h) , for which the power allocation is not adapted to the channel and observation states.

Proposition 8. *Policies in Π^{sub1} are \bar{D} -feasible if the following conditions hold:*

$$\sum_q \sum_h \Pr(q) \Pr(h) \mathbb{E} [f^q (D^{q,h}, \alpha^{q,h} E_k)] < \sum_q \sum_h \Pr(q) \Pr(h) \mathbb{E} [g^h ((1 - \alpha^{q,h}) E_k)], \quad (13)$$

$$\text{and } \sum_q \sum_h \Pr(q) \Pr(h) D^{q,h} \leq \bar{D}, \quad (14)$$

where the expectation \mathbb{E} in (13) is over the energy harvesting process E_k .

Remark 9. In general, the set of policies Π^{sub1} is not guaranteed to be a distortion-optimal energy-neutral class, since the necessary conditions of Proposition 3 could hold where the sufficient conditions of Proposition 8 do not. This is also confirmed via numerical simulations in Sec. III-D. However, for constant observation and channel states, i.e., $H_k = h_0$ and $Q_k = q_0$ for all k , and for f and g linear in $T_{s,k}$ and $T_{t,k}$, respectively, the class Π^{sub1} is distortion-optimal energy-neutral. In fact, under these assumptions, the sufficient condition (13) becomes $f^{q_0} (D^{q_0,h_0}, \alpha^{q_0,h_0} \mathbb{E} [E_k]) < g^{h_0} ((1 - \alpha^{q_0,h_0}) \mathbb{E} [E_k])$. Defining $T_s^q = \alpha^{q_0,h_0} \mathbb{E} [E_k]$, $T_t^h = (1 - \alpha^{q_0,h_0}) \mathbb{E} [E_k]$ and $D^q = D^{q_0,h_0}$, conditions (13)-(14) correspond to (8). Thus, the class of policies Π^{sub1} is distortion-optimal energy-neutral.

Proof: Due to the stationarity and ergodicity of the processes Q_k and H_k , the parameters D_k , $T_{s,k}$ and $T_{t,k}$, are also stationary ergodic and (13) is a sufficient condition for the stability of the queue \tilde{X}_k [20, Ch.3]. Since $\lim_{n \rightarrow \infty} \frac{1}{n} \sum_{k=1}^n D_k = \sum_q \sum_h \Pr(q) \Pr(h) D^{q,h}$, for $D^{q,h}$ such that $\sum_q \sum_h \Pr(q) \Pr(h) D^{q,h} \leq \bar{D}$ the class of policies Π^{sub1} satisfies the constraint (7). ■

The greedy policies introduced above do not make use of the energy buffer at all, whereas the distortion-optimal energy-neutral class of policies Π^{do} does. For comparison purposes, it is interesting to consider hybrid policies that differ from those in Π^{do} as the energy buffer is used only either for compression or for transmission. The first policies Π^{hyb1} require an energy buffer for the channel encoder only in order to adapt the transmission power to the channel state, i.e., $T_{t,k} = \min[\alpha \tilde{E}_k, T_t^h]$ for $H_k = h$. The energy allocated to the source encoder is instead independent of the observation state, i.e., $T_{s,k} = (1 - \alpha) E_k$. Viceversa, the second policies Π^{hyb2} are adapted to the observation state instead of the channel state, and require an energy buffer only for the source encoder, i.e., $T_{s,k} = \min[(1 - \alpha) \tilde{E}_k, T_s^q]$ for $Q_k = q$ and $T_{t,k} = \alpha E_k$.

Proposition 10. *Policies in Π^{hyb1} are \bar{D} -feasible if conditions (8)-right and (11)-right hold, along with $\sum_q \Pr(q) \mathbb{E}[f^q(D^q, (1 - \alpha) E_k)] < \sum_h \Pr(h) g^h(T_t^h)$. Similarly, policies Π^{hyb2} are \bar{D} -feasible if conditions (8)-right and (11)-left hold, along with $\sum_q \Pr(q) f^q(D^q, T_s^q) < \sum_h \Pr(h) \mathbb{E}[g^h(\alpha E_k)]$.*

Proof: Proof follows from the proofs of Propositions 5 and 8. ■

C. Analog Transmission

In this section, we consider for performance comparison an alternative class of strategies in which the sampled source is transmitted directly via analog modulation (see, e.g., [15]). In other words, a block of source samples is scaled and transmitted in one slot, so as to consume $T_{t,k}$ transmission energy per channel use during slot k . Energy $T_{t,k}$ is selected as $T_{t,k} = \min[\tilde{E}_k - \epsilon, T_t^{q,h}]$ for $Q_k = q$ and $H_k = h$ for given parameters $T_t^{q,h} \geq 0$, so that it depends on the current source and channel states. If the bandwidth ratio $b = N/M$ is larger than one, i.e., there are more channel uses than source samples, the extra $N - M$ source samples are unused. Instead,

if $b < 1$, then a fraction $1 - b$ of source samples is not transmitted. Notice that this class of strategies does not fall in the category depicted in Fig. 1 and discussed above, since it does not have separate encoders.

We assume the observation model of Example 1 with an i.i.d. source $U_{k,i} \sim \mathcal{N}(0, d_{max})$ and an AWGN channel with SNR H_k . For a bandwidth ratio of b , it is not difficult to obtain that the MMSE at the receiver is given by

$$d_{mmse}(T_{t,k}, Q_k, H_k) = \begin{cases} \left(\frac{bT_{t,k}Q_kH_k}{bT_{t,k}Q_k + Q_k + 1} + \frac{1}{d_{max}} \right)^{-1} & \text{for } b \geq 1 \\ b \times \left(\frac{T_{t,k}Q_kH_k}{T_{t,k}Q_k + Q_k + 1} + \frac{1}{d_{max}} \right)^{-1} + (1 - b)d_{max} & \text{for } b < 1 \end{cases}. \quad (15)$$

Notice that, for fairness, the average energy used for the transmission of one sample is $bT_{t,k}$ if $b \geq 1$. Also, notice that if $b < 1$, the maximum distortion d_{max} is accrued on the fraction $(1 - b)$ of samples that are not transmitted. For simplicity, we assume that analog transmission has negligible power spent for source acquisition, i.e., $T_{s,k} = 0$, though this is not entirely correct given that even in this case there is a need for sensing, sampling and analog-to-digital conversions. Nonetheless, these power consumption terms are also neglected in Examples 1 and 2. Under this assumption, we have the following.

Proposition 11. *Analog transmission satisfies the distortion constraint (7) if the following conditions are satisfied:*

$$\sum_q \sum_h \Pr(q) \Pr(h) d_{mmse}(T_t^{q,h}, Q = q, H = h) \leq \bar{D}, \quad \text{and} \quad \sum_q \sum_h \Pr(q) \Pr(h) T_t^{q,h} \leq \mathbb{E}[E_k]. \quad (16)$$

Proof: Follows similar to Proposition 6. ■

We can also consider a greedy policy $T_{t,k} = E_k$, for which the power allocation is not adapted to the channel and observation states and energy storage is not required.

D. Numerical Results

In this section we compare numerically the performance of the optimal and suboptimal source-channel coding policies presented so far, along with analog transmission strategies.

Consider first a scenario where the observation and channel states are constant, i.e., $Q_k = q$ and $H_k = h$ for all k . The energy arrival E_k has mean 1 Joule/channel use and uniform pdf between 0 and 2 Joule/channel use. We consider model (3) with $T_s^{max} = 1$ Joule/source sample, efficiency parameters $\zeta = 1$ and $\eta = 1.5$ for the source encoder, and the complex AWGN channel Shannon capacity $g^h(T_{t,k}) = N \times \log(1 + hT_{t,k})$. In Fig. 2, we identify the values of source and channel SNRs (q, h) for which different policies are able to stabilize the data queue and guarantee average distortion $\bar{D} = 0.8$. We refer to these regions as “achievable regions”. Achievable regions are given in Fig. 2 by the area above the corresponding lines. We use standard tools of convex optimization for their numerical evaluation.

In Fig. 2-b, we can further observe that the achievable regions of the distortion-optimal energy-neutral class (10) are significantly larger than those of the greedy policies (12) due to possibility to store energy and thus allocate resources more effectively. Moreover, by considering also the hybrid policies Π^{hyb1} and Π^{hyb2} , we can see that most of the gains are obtained, in this example, by exploiting the energy buffer in order to allocate energy over time to the source encoder, whereas the gains accrued by using the battery for data transmission are less significant. This is observed by noticing that the achievable region of the class Π^{do} is close to that obtained by hybrid policies Π^{hyb2} , but much larger than that obtained by hybrid class of policies Π^{hyb1} . The relative comparison between the two hybrid policies, and thus between the use of the battery for source or channel encoding, depends on the functions $f(D_k, T_{s,k}, Q_k)$ and $g(H_k, T_{t,k})$. For instance, setting a lower T_s^{max} would change the presented results by penalizing more the strategies that are not using the energy buffer for channel transmission.

We now consider the performance of analog transmission. As it is well known from rate-distortion theory [18], for bandwidth ratio $b = 1$ and the given (Gaussian) source and channel models, analog transmission is rate-distortion optimal. Separate source-channel coding is also optimal (for any $b > 0$) if compression is assumed not to consume any energy. Here, instead, the achievable region of analog transmission is expected to be larger than that of strategies that employ separate source-channel coding, as source encoding energy costs are taken into account in the given model (3). This is confirmed by Fig. 2-b. However, for sufficiently larger or smaller

bandwidth ratios, the extra energy spent for compression is not enough to overcome the rate-distortion gains attained by separate source-channel coding versus analog transmission. This is apparent from Fig. 2-a and Fig. 2-c where the analog transmission is outperformed.

We now consider a scenario where source and channel states are not constant but vary with two possible states, namely $\mathcal{Q} = \{10^{-1}, 10^2\}$ (Fig. 3-a) or $\mathcal{Q} = \{10^{-0.2}, 1\}$ (Fig. 3-b) for source SNR, and $\mathcal{H} = (10^{-1}, 10^2)$ (Fig. 3-a) or $\mathcal{H} = (3.5, 7)$ (Fig. 3-b) for channel SNR. The worst-case observation SNR (e.g., $Q_k = 10^{-1}$ for Fig. 3-a) and the worst-case channel SNR (e.g., $H_k = 10^{-1}$ for Fig. 3-a) have probabilities p_w^q and p_w^h , respectively. In Fig. 3, achievability regions are the sets of probability values (p_w^q, p_w^h) for which different policies guarantee queue stability and average distortion $\bar{D} = 0.8$. In particular, the regions are identified as the area below the corresponding curves. The results emphasize the importance of jointly adapting the resource allocation to both source and channel states in case of a greedy policy that does not employ the battery. This is seen by comparing the performance of the greedy schemes Π^{sub1} (12), which adapts the policy to the current states, and Π^{sub2} , which does not. Moreover, comparing Fig. 3-b with Fig. 3-a, it is seen that the better “worst-case” state allows the distortion-optimal energy-neutral policy to satisfy stability and average distortion constraints for larger values of the probabilities (p_w^q, p_w^h) . On the contrary, the greedy policies suffer from the worse “best-case” state $(q, h) = (1, 7)$ of Fig. 3-b, as this corresponds to operating critically close to the border of their achievable regions (see the achievable region of Π^{sub1} in Fig. 2-b).

IV. DELAY-DISTORTION OPTIMIZATION

The stability criterion considered in Sec. III does not provide any guarantee on the delay experienced by the reconstruction of the source in a certain time-slot. In some applications, instead, one may be willing to trade distortion for a shorter delay. In this section, we address such requirement by looking for policies that minimize a weighted sum of distortion and delay. In particular, we propose to minimize the *expected total discounted cost* [22]

$$\lim_{n \rightarrow \infty} \frac{1}{n} \sum_{k=0}^n \lambda^k \left[\gamma D_k + (1 - \gamma) \tilde{X}_k \right], \quad (17)$$

where $0 \leq \lambda < 1$ is the discount factor and $0 \leq \gamma \leq 1$. The latter parameter weights the importance of distortion versus delay in the optimization criterion. Notice that if $\gamma = 0$ then one minimizes the average length of the data queue, which, by Little's theorem, is the same as minimizing the average delay.

In order to tackle the minimization of (17) over the policies π defined in Sec. III, we assume that: (i) The data and energy buffers are finite; (ii) The set of possible decisions $\pi_k := \{D_k, T_{s,k}, T_{t,k}\}$ is discrete; (iii) The energy arrival E_k takes values in a discrete and finite set; (iv) The sets of values assumed by rates $f(D_k, T_{s,k}, Q_k)$, $g(H_k, T_{t,k})$, and by the queue length \tilde{X}_k are discrete. Following standard theory [22, Ch. 6], these assumptions entail that the optimal policy is deterministic and stationary (Markovian). In other words, $\{D_k, T_{s,k}, T_{t,k}\}$ are function of the present state $S_k = \{\tilde{E}_k, \tilde{X}_k, Q_k, H_k\}$ only. Therefore, the solution can be found via value iteration [22]. Notice that, due to (i), data buffer overflow may happen, in which case the compression bits are lost and a maximum distortion d_{max} is accrued for the current slot.

While in general the optimal policy allocates resources to source and channel encoder through parameters $T_t^{q,h}$ as a function of both source $Q_k = q$ and channel $H_k = h$ states, the class of policies Π^{do} (10) performs such allocation independently for source and channel encoders. For comparison purposes, we evaluate also the performance in terms of criterion (17) of a class of policies that optimize separately the source encoder parameters $\{D_k, T_{s,k}\}$ as a function of Q_k , and the channel encoder parameter $\{T_{t,k}\}$ as a function of H_k . As for the distortion-optimal energy-neutral class of policies, the energy resources are split between the encoders, such that the source encoder makes use of a fraction α of the energy, whereas the rest is utilized by the channel encoder. Specifically, the energy-buffer is divided into two buffers, that are used independently by the encoders: the source encoder buffer is charged by αE_k , while the channel encoder buffer absorbs the remaining quantity of energy arrival $(1 - \alpha) E_k$. The source encoder policy $\{D_k, T_{s,k}\}$ is optimized via value iteration with respect to the criterion (17) by assuming a constant transmission rate $g(H_k, T_{t,k}) = \bar{g}$. On the other hand, the channel encoder policy $\{T_{t,k}\}$ is optimized via value iteration with respect to criterion (17) with weight $\gamma = 0$ (since it cannot optimize its policy with respect to the distortion), by assuming a constant source rate

$f(D_k, T_{s,k}, Q_k) = \bar{f}$. The best “separable” policy is finally obtained by selecting the values $(\alpha, \bar{g}, \bar{f})$ that achieve the best delay-distortion trade-off.

A. Numerical Results

In this section we compute numerically the trade-off between delay and distortion by minimizing (17) for different values of γ . Specifically, for each γ , we evaluate the average delay, which is measured by the average data queue length by Little’s law, and average distortion. The optimal policies are computed via value iteration [22] and so are the suboptimal policies corresponding to separate optimization of source and channel encoders.

In Fig. 4 the delay-distortion trade-off is shown both for the optimal policies and for the “separable” ones discussed above. The discount factor is $\lambda = 0.5$. The compression model is (5), with minimum required energy per sample $\nu = 0.1$ Joule/sample and bandwidth ratio $b = 1$. The quantities of interest are discretized as follows: $\tilde{X}_k \in \{0, \dots, 5\}$ is expressed in multiples of the codeword length $M = N$; The energy buffer size is $\tilde{E}_k - E_k \in \{0, 1, 2\}$ and $E_k \in \{1, 2\}$ with p_w^e being the probability that $E_k = 1$ (worst case); The source correlation values are $\mathcal{Q} = \{0.1, 0.5\}$ and channel SNR values are $\mathcal{H} = (0.5, 10)$, with probabilities p_w^q and p_w^h for $Q_k = 0.1$ and $H_k = 0.5$ (worst cases); The distortion takes values as $D_k \in \{0.1, 0.55, 1\}$ and $d_{max} = 1$; The source-encoder rate $f(D_k, T_{s,k}, Q_k)/M$ is rounded to the smallest following integer, while the channel-encoder rate $g(H_k, T_{t,k})/N$ is rounded to the largest previous integer.

In Fig. 4 we observe that the optimal policies obtain a remarkably better delay-distortion trade-off compared to the separable policies, both for low and large worst-case probabilities, i.e., $p_w^e = p_w^q = p_w^h = p_w = 0.1$ and 0.9 . This demonstrates the importance of a joint resource allocation over the encoders whenever the delay is also of interest. Note that for increasing average buffer length (delay), since the buffer size is finite, it becomes more crucial to adopt a joint resource allocation. This is because the separate approach is not able to prevent buffer overflow as the source encoder operates without channel state information.

V. MULTIPLE ACCESS

In this section, we briefly discuss an extension of the analysis to a scenario in which L sensors access a single access point employing TDMA. Random access protocols will be considered in future work (an analysis with exogenous bit arrivals, and thus no source encoder, can be found in [13]). Each sensor is modeled as described in Sec. II, and we assume that observation qualities $\mathbf{Q}_k = [Q_{1,k}, \dots, Q_{L,k}]^T \in \mathcal{Q}^L = [\mathcal{Q}_1, \dots, \mathcal{Q}_L]$, channel qualities $\mathbf{H}_k = [H_{1,k}, \dots, H_{L,k}]^T \in \mathcal{H}^L = [\mathcal{H}_1, \dots, \mathcal{H}_L]$, and the energy arrivals $\mathbf{E}_k \in \mathbb{R}_+^L$ are jointly stationary and ergodic. We tackle the problem of designing policies defined, extending Sec. II-B, as the tuple $v := \{\tau_k, \boldsymbol{\pi}_k\}_{k \geq 1}$ that consists of the scheduling policy $\tau_k \in \{1, \dots, L\}$, which reserves the slot k to one sensor $l \in \{1, \dots, L\}$, and of the joint resource allocation policy $\boldsymbol{\pi}_k = [\pi_{1,k}, \dots, \pi_{L,k}]^T$, where each entry $\pi_{l,k}$ is defined as in Sec. II-B and corresponds to the distortion and energy allocation for the l^{th} sensor. Recall that $\{\tau_k, \boldsymbol{\pi}_k\}$ generally depends on the whole history of past and current states (see Sec. II-B). Note that time-slot k is exclusively assigned to sensor l , i.e., $T_{t,l,k} = 0$ if $\tau_k \neq l$. We define a policy v as $\bar{\mathbf{D}}$ -feasible if it guarantees the stability of all data queues and average distortion constraints $\lim_{n \rightarrow \infty} \frac{1}{n} \sum_{k=1}^n D_{k,l} \leq \bar{D}_l$, collected for notational convenience in vector $\bar{\mathbf{D}} = [\bar{D}_1, \dots, \bar{D}_L]^T$. As for the single sensor scenario, we are interested in finding a distortion-optimal energy-neutral class of policies $\Upsilon^{do} \subseteq \Upsilon$, i.e., a subset of all possible scheduling policies Υ that contains at least one $\bar{\mathbf{D}}$ -feasible policy.

In the following we state a necessary condition for the existence of a $\bar{\mathbf{D}}$ -feasible policy v .

Proposition 12. *For a set of distortion constraints $\bar{\mathbf{D}}$, a necessary condition for the existence of a $\bar{\mathbf{D}}$ -feasible policy v is the existence of the set of parameters $D_l^{q_l} \geq 0$, $T_{s,l}^{q_l} \geq 0$ for $q_l \in \mathcal{Q}_l$, $T_{t,l}^{h_l} \geq 0$ for $h_l \in \mathcal{H}_l$, $0 < \beta_l^{\mathbf{h}} < 1$, for $\mathbf{h} \in \mathcal{H}^L$, and $0 < \alpha_l < 1$, such that*

$$\sum_{q_l \in \mathcal{Q}_l} \Pr(q_l) f^{q_l}(D_l^{q_l}, T_{s,l}^{q_l}) < \sum_{h_l \in \mathcal{H}_l} \left(\sum_{\mathbf{h} \in \mathcal{H}^L: \mathbf{h}(l)=h_l} \Pr(\mathbf{h}) \beta_l^{\mathbf{h}} \right) g^{h_l}(T_{t,l}^{h_l}), \quad (18)$$

$$\sum_{l=1}^L \beta_l^{\mathbf{h}} = 1 \text{ for all } \mathbf{h}, \quad \sum_{q_l \in \mathcal{Q}_l} \Pr(q_l) D_l^{q_l} \leq \bar{D}_l \quad (19)$$

$$\sum_{q_l \in \mathcal{Q}_l} \Pr(q_l) T_{s,l}^{q_l} \leq (1 - \alpha_l) \mathbb{E}[E_{l,k}], \quad \text{and} \quad \sum_{h_l \in \mathcal{H}_l} \left(\sum_{\mathbf{h} \in \mathcal{H}^L: \mathbf{h}(l) = h_l} \Pr(\mathbf{h}) \beta_l^{\mathbf{h}} \right) T_{t,l}^{h_l} \leq \alpha_l \mathbb{E}[E_{l,k}]. \quad (20)$$

Remark 13. The interpretation of Proposition 12 is similar to the one of Proposition 3 given in Remark 4. The additional parameters $\beta_l^{\mathbf{h}}$ can be interpreted as the fraction of the subset of time slots with joint channel state equal to \mathbf{h} for which the sensor l is scheduled.

Proof: A necessary condition for the stability of the queue of the l^{th} sensor is $\mathbb{E}_v[f(D_{l,k}, T_{s,l,k}, Q_{l,k})] < \mathbb{E}_v[g(H_{l,k}, T_{t,l,k})]$ [20]. Since we have $T_{t,l,k} = 0$ for $\tau_k \neq l$, on the right of the inequality we obtain $\mathbb{E}_v[g(H_{l,k}, T_{t,l,k})] = \sum_{h_l \in \mathcal{H}_l} \sum_{\mathbf{h} \in \mathcal{H}^L: \mathbf{h}(l) = h_l} \Pr_v(\tau_k = l, \mathbf{H}_k = \mathbf{h}) \times \mathbb{E}_v[g(H_{l,k}, T_{t,l,k}) | \tau_k = l, H_{l,k} = h_l]$. Then, using Jensen inequality at both sides, the stability condition becomes

$$\begin{aligned} \sum_{q_l \in \mathcal{Q}_l} \Pr(q_l) f^{q_l}(\mathbb{E}_v[D_{l,k} | Q_{l,k} = q_l], \mathbb{E}_v[T_{s,l,k} | Q_{l,k} = q_l]) &\leq \mathbb{E}_v[f(D_{l,k}, T_{s,l,k}, Q_{l,k})] < \\ &< \sum_{h_l \in \mathcal{H}_l} \sum_{\mathbf{h} \in \mathcal{H}^L: \mathbf{h}(l) = h_l} \Pr_v(\tau_k = l, \mathbf{H}_k = \mathbf{h}) \mathbb{E}_v[g(H_{l,k}, T_{t,l,k}) | \tau_k = l, H_{l,k} = h_l] \leq (21) \\ &\leq \sum_{h_l \in \mathcal{H}_l} \sum_{\mathbf{h} \in \mathcal{H}^L: \mathbf{h}(l) = h_l} \Pr(\mathbf{h}) \Pr_v(\tau_k = l | \mathbf{H}_k = \mathbf{h}) g^{h_l}(\mathbb{E}_v[T_{t,l,k} | \tau_k = l, H_{l,k} = h_l]). \end{aligned}$$

Finally, condition (18) is obtained for $D_l^{q_l} = \mathbb{E}_v[D_{l,k} | Q_{l,k} = q_l]$, $T_{s,l}^{q_l} = \mathbb{E}_v[T_{s,l,k} | Q_{l,k} = q_l]$, $\beta_l^{\mathbf{h}} = \Pr_v(\tau_k = l | \mathbf{H}_k = \mathbf{h})$, and $T_{t,l}^{h_l} = \mathbb{E}_v[T_{t,l,k} | \tau_k = l, H_{l,k} = h_l]$. Note that (19)-left follows immediately from this definition. As for conditions (20), recall that, from (1), we must have $\frac{1}{K} \sum_{k=1}^K (T_{s,l,k} + T_{t,l,k}) \leq \frac{1}{K} \sum_{k=1}^K E_{l,k} + \frac{\tilde{E}_{l,0}}{K}$, for $K \geq 1$, and the initial state of the energy buffer $\tilde{E}_{l,0}$. Thus, for a stationary ergodic policy v , we get

$$\mathbb{E}_v[T_{s,l,k}] + \sum_{h_l \in \mathcal{H}_l} \sum_{\mathbf{h} \in \mathcal{H}^L: \mathbf{h}(l) = h_l} \Pr(\mathbf{h}) \Pr_v(\tau_k = l | \mathbf{H}_k = \mathbf{h}) \mathbb{E}_v[T_{t,l,k} | \tau_k = l, H_{l,k} = h_l] \leq \mathbb{E}[E_{l,k}], \quad (22)$$

where $\frac{1}{K} \sum_{k=1}^K T_{s,l,k} \rightarrow \mathbb{E}_v[T_{s,l,k}]$, $\frac{1}{K} \sum_{k=1}^K E_{l,k} + \frac{\tilde{E}_{l,0}}{K} \rightarrow \mathbb{E}[E_{l,k}]$, and $\frac{1}{K} \sum_{k=1}^K T_{t,l,k} \rightarrow \mathbb{E}_v[T_{t,l,k}] = \sum_{h_l \in \mathcal{H}_l} \sum_{\mathbf{h} \in \mathcal{H}^L: \mathbf{h}(l) = h_l} \Pr(\mathbf{h}) \Pr_v(\tau_k = l | \mathbf{H}_k = \mathbf{h}) \mathbb{E}_v[T_{t,l,k} | \tau_k = l, H_{l,k} = h_l]$. Given the above definition of $\beta_l^{\mathbf{h}}$, $T_{s,l}^{q_l}$ and $T_{t,l}^{h_l}$, inequality (22) becomes $\sum_{q_l \in \mathcal{Q}_l} \Pr(q_l) T_{s,l}^{q_l} +$

$+ \sum_{h_l \in \mathcal{H}_l} \sum_{\mathbf{h} \in \mathcal{H}^L: \mathbf{h}(l)=h_l} \Pr(\mathbf{h}) \beta_l^{\mathbf{h}} T_{t,l}^{h_l} \leq \mathbb{E}[E_{l,k}]$, proving (20), where $\alpha_l = \mathbb{E}_v[T_{t,l,k}] / \mathbb{E}[E_{l,k}]$. To conclude, condition (19)-right follows from the distortion constraints similar to Proposition 3. \blacksquare

In order to define a distortion-optimal energy-neutral class of policies, Proposition 12 suggests to consider a class of scheduling policies Υ^{do} in which scheduling is done opportunistically based on the channel states \mathbf{h} of all sensors according to a probability distribution $\beta_l^{\mathbf{h}}$: if the channels are equal to \mathbf{h} , then sensor l is selected with probability $\beta_l^{\mathbf{h}}$. Notice that scheduling is independent of the observation qualities in a given slot. Moreover, energy and distortion allocations at each sensors are similar to the policies π_k in the class Π^{do} (10), and, thus, in particular perform separate resource allocation over source and channel encoders. It can be shown that, similar to Proposition 5, the so defined class of policies Υ^{do} is distortion-optimal energy-neutral in Υ .

A. Numerical Results

In this section we assess numerically the performance of the distortion-optimal energy-neutral class of scheduling policies Υ^{do} . For comparison purposes, we introduce the suboptimal class of policies Υ^{sub} , that schedules each sensor according to a fixed probability $\beta_l = \Pr(\tau_k = l)$, independently of the current channel conditions. We consider $L = 2$ sensors, which are modeled as in Fig. 3-b (see Sec. III-D). For sensor 2, the probabilities $(p_w^{q_2}, p_w^{h_2})$ of the worst observation and channel states are fixed, whereas, for sensor 1, $(p_w^{q_1}, p_w^{h_1})$ are varied. Fig. 5 shows the corresponding achievability region, defined, as in Sec. III-D, as the set of $(p_w^{q_1}, p_w^{h_1})$ for which the given policy is able to stabilize the data queues and guarantee the given average distortions. For further comparison, Fig. 5 also shows the outer bound to the achievability region given by the case where only sensor 1 is present. The numerical results confirm that the achievability region of the optimal class of strategies Υ^{do} (dashed lines) is larger than that of the suboptimal class Υ^{sub} (dot-dashed lines). Moreover, note that the achievability regions shrink if the worst case states probabilities $(p_w^{q_2}, p_w^{h_2})$ of the second sensor get larger, since sensor 2 requires more transmission resources to compensate for both the worst observation and channel conditions. It is further interesting to observe that, for $p_w^{h_2} = 0.1$ and $p_w^{h_1} = 1$, the achievability regions of Υ^{do}

and Υ^{sub} , in terms of $p_w^{q_1}$, are practically the same (circle marker). This is due to the fact that the variations of channels \mathbf{H}_k are not large enough to enables gains by adapting parameters β_l^h .

VI. CONCLUSIONS

We studied energy management for a system consisting of a single sensor whose task is that of reporting the measure of a phenomenon to a receiver. The main problem is that of allocating energy between the source and the channel encoders based on the current amount of available energy, state of the data queue, quality of the measurement and of the wireless channel. We first look for a distortion-optimal energy-neutral subset of all policies, that contains at least one policy able to stabilize the data queue and to satisfy a maximum average distortion constraint. We found that optimal policies according to this criterion operate a separate energy allocation of source and channel encoder. Instead, we showed that a joint energy management over source and channel encoder is required to achieve the desired trade-off between delay and distortion. Finally, we considered a system with multiple sensors and obtained TDMA scheduling policies that guarantee the stability of all data queues, whenever the distortion constraints are feasible. Overall, our results, which also include further comparisons with a number of suboptimal policies, shed light on the challenges and design issues that characterize modern cyber-physical systems.

ACKNOWLEDGMENTS

The authors would like to thank Prof. Petar Popovski of Aalborg University for suggesting the analysis of the analog transmission scheme defined in Sec. III-C.

REFERENCES

- [1] "It's a smart world," *The Economist*, Nov. 6 2010.
- [2] J. A. Paradiso and T. Starner, "Energy scavenging for mobile and wireless electronics," *IEEE Pervasive Computing*, vol. 4, pp. 18–27, Jan.-Mar. 2005.
- [3] R. Rajkumar, L. Insup, S. Lui, and J. Stankovic, "Cyber-physical systems: The next computing revolution," in *Proc. ACM/IEEE Design Automation Conf. (DAC)*, June 13-18 2010.
- [4] A. Kansal, J. Hsu, S. Zahedi, and M. B. Srivastava, "Power management in energy harvesting sensor networks," *ACM Trans. Embed. Comput. Syst.*, vol. 6, Sep. 2007.

- [5] Z. He and D. Wu, "Resource allocation and performance analysis of wireless video sensors," *IEEE Trans. on Circ. and Syst. for Video Technol.*, vol. 16, pp. 590 – 599, May 2006.
- [6] X. Lu, E. Erkip, Y. Wang, and D. Goodman, "Power efficient multimedia communication over wireless channels," *IEEE J. on Sel. Areas in Commun.*, vol. 21, pp. 1738 – 1751, Dec. 2003.
- [7] V. Sharma, U. Mukherji, V. Joseph, and S. Gupta, "Optimal energy management policies for energy harvesting sensor nodes," *IEEE Trans. Wireless Commun.*, vol. 9, pp. 1326–1336, Apr. 2010.
- [8] O. Ozel and S. Ulukus, "Information-theoretic analysis of an energy harvesting communication system," in *Proc. IEEE 21st Int. Symp. on Personal, Indoor and Mobile Radio Commun. (PIMRC)*, Sep. 2010.
- [9] R. Rajesh and V. Sharma, "Capacity of fading Gaussian channel with an energy harvesting sensor node," Submitted, available at <http://arxiv.org/abs/1010.5416>.
- [10] Z. Mao, C. Koksas, and N. Shroff, "Resource allocation in sensor networks with renewable energy," in *Proc. IEEE 19th Int. Conf. on Computer Commun. and Netw. (ICCCN)*, Aug. 2010.
- [11] J. Yang and S. Ulukus, "Optimal packet scheduling in an energy harvesting communication system," Submitted, available at <http://arxiv.org/pdf/1010.1295v1>.
- [12] K. Tutuncuoglu and A. Yener, "Optimum transmission policies for battery limited energy harvesting nodes," Submitted, available at <http://arxiv.org/abs/arXiv:1010.6280>.
- [13] V. Sharma, U. Mukherji, and V. Joseph, "Efficient energy management policies for networks with energy harvesting sensor nodes," in *Proc. 46th Annual Allerton Conf. on Commun., Control, and Computing*, Sep. 2008.
- [14] R. Srivastava and C. E. Koksa, "Basic tradeoffs for energy management in rechargeable sensor networks," Submitted, available at <http://arxiv.org/abs/1009.0569v1>.
- [15] B. R. M. Gastpar and M. Vetterli, "To code, or not to code: lossy source–channel communication revisited," *IEEE Trans. Inf. Theory*, vol. 49, pp. 1147–1158, May 2003.
- [16] T. M. Cover and J. A. Thomas, *Elements of information theory*. Wiley.
- [17] T. Burd and R. Broderon, "Processor design for portable systems," *J. VLSI Signal Process.*, vol. 13, pp. 203–222, Aug. 1996.
- [18] T. Berger, *Rate distortion theory*. Prentice Hall, 1971.
- [19] A. Gersho and R. M. Gray, *Vector quantization and signal compression*. Kluwer, 1992.
- [20] A. A. Borovkov, *Stochastic processes in queueing theory*. Springer-Verlag, 1976.
- [21] L. Tassiulas and A. Ephremides, "Stability properties of constrained queueing systems and scheduling policies for maximum throughput in multihop radio networks," *IEEE Trans. Autom. Control*, vol. 37, pp. 1936–1948, Dec. 1992.
- [22] M. L. Puterman, *Markov decision processes*. Wiley, 1994.




Figure 1. An energy-harvesting sensor composed of a cascade of a source and a channel encoder powered by a resource manager that allocates the energy available in the buffer (e.g., battery or capacitor).

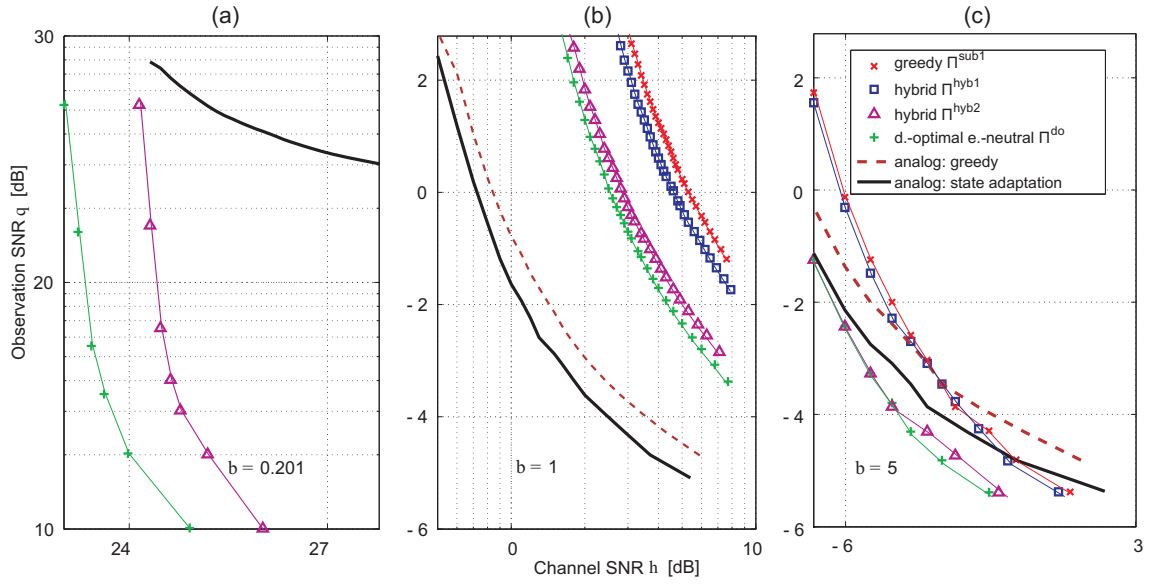


Figure 2. Achievable regions (regions above the curves) for the digital policies (lines with markers) and for the analog policies (only lines): (a) bandwidth ratio $b = 0.201$; (b) bandwidth ratio $b = 1$; (c) bandwidth ratio $b = 5$ ($E_k \sim \mathcal{U}(0, 2)$, with mean 1 Joule/channel use; $\bar{D} = 0.8$; compression model (3), with $T_s^{\max} = 1$ Joule/source sample, $\zeta = 1$, $\eta = 1.5$, maximum distortion $d_{\max}=1$).

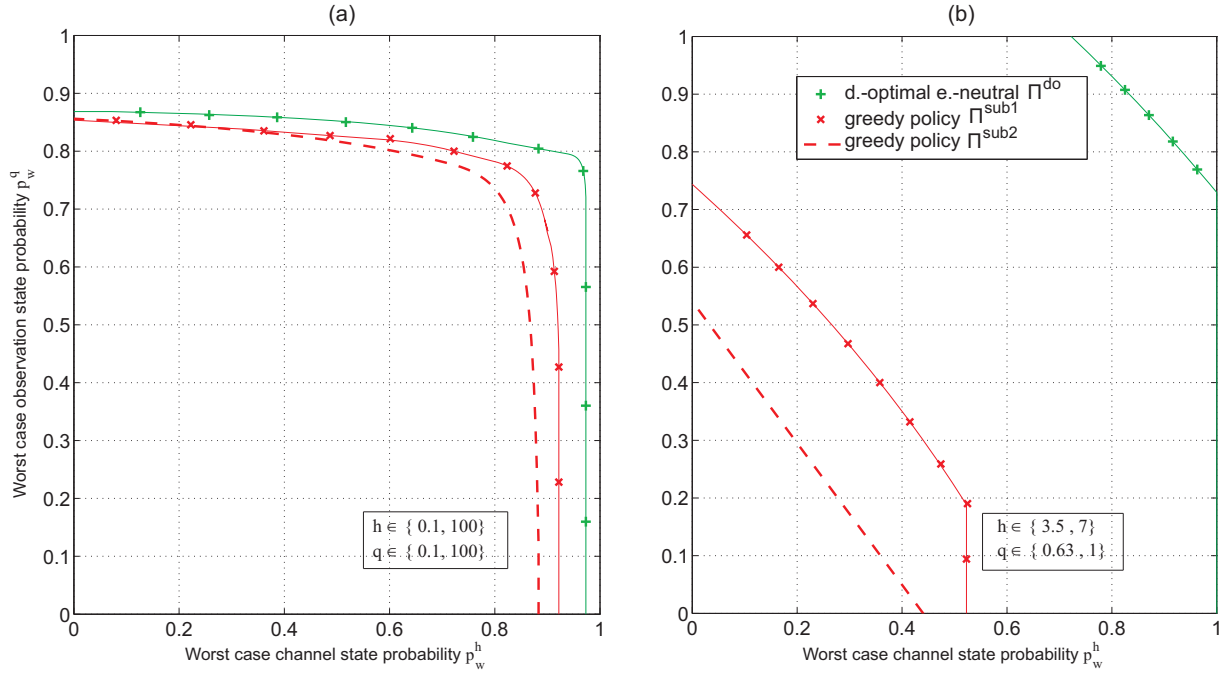


Figure 3. Achievable regions (regions below the curves) of the digital policies Π^{do} (10), Π^{sub1} (12), and Π^{sub2} , with two channel and observation SNR states, respectively. ($E_k \sim \mathcal{U}(0, 2)$, with mean 1 Joule/channel use; $\bar{D} = 0.8$; compression model (3), with $T_s^{max} = 1$ Joule/source sample, $\zeta = 1$, $\eta = 1.5$, bandwidth ratio $b = 1$, maximum distortion $d_{max}=1$).




Figure 4. Delay-distortion trade-off, where average delay is proportional to the depicted average data queue length (maximum data-buffer length: 5 codeword lengths; maximum distortion: $d_{max}=1$; discount factor: $\lambda = 0.5$; compression model (5), with minimum required energy per sample $\nu = 0.1$ Joule/sample; bandwidth ratio: $b = 1$; queue length: $\tilde{X}_k \in \{0, \dots, 5\}$ expressed in multiples of M ; energy buffer size: $\tilde{E}_k - E_k \in \{0, 1, 2\}$; energy arrival: $E_k \in \{1, 2\}$ Joule/sample; source correlation values: $\mathcal{Q} = \{0.1, 0.5\}$; channel SNR values: $\mathcal{H} = (0.5, 10)$; distortion values: $D_k \in \{0.1, 0.55, 1\}$).




Figure 5. Achievable regions (regions below the curves) of the scheduling policies Υ^{do} (dashed lines) and Υ^{sub} (dot-dashed lines). The solid line corresponds to the optimal policies when the second sensor $l = 2$ is not present ($E_{l,k} \sim \mathcal{U}(0, 2)$, with mean 1 Joule/channel use; $\bar{D}_l = 0.8$; compression model (3), with, for both sensors, $T_s^{max} = 1$ Joule/source sample, $\zeta = 1$, $\eta = 1.5$, bandwidth ratio $b = 1$, maximum distortion $d_{max}=1$).